AMD
Smarter Choice

Next-Generation Mobile Computing:
Balancing Performance and Power Efficiency

HOT CHIPS 19
Jonathan Owen, AMD

Agenda

AMD
Smarter Choice

- The mobile computing evolution
- The "Griffin" architecture
- Memory enhancements
- Power management
- Thermal management

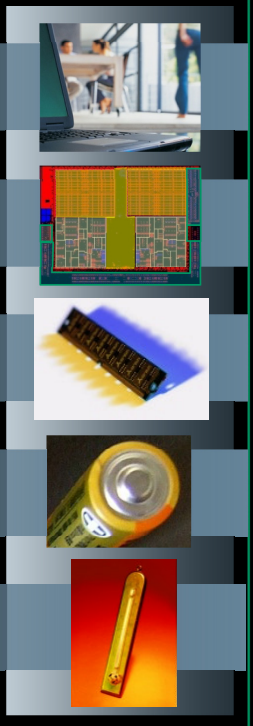

The Evolution of the Notebook PC

Customers increasingly demand more for less

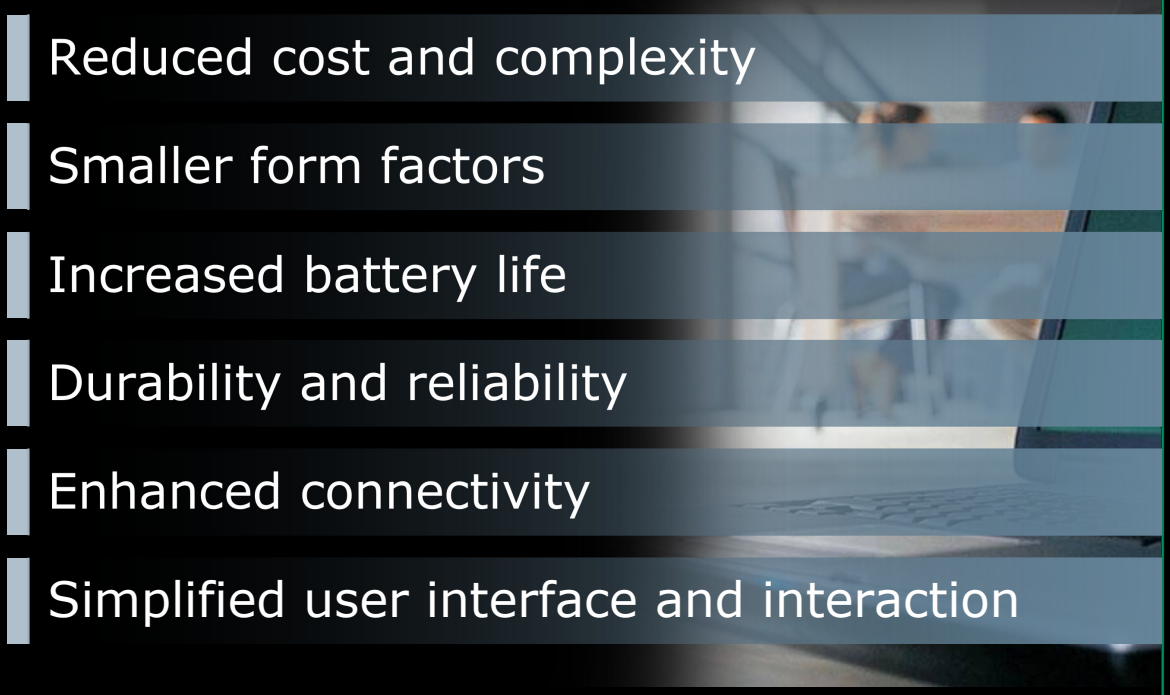

- Reduced cost and complexity
- Smaller form factors
- Increased battery life
- Durability and reliability
- Enhanced connectivity
- Simplified user interface and interaction

Architectural Challenges for Northbridges in Mobile Platforms

- UMA configurations present increasing challenges:
  - Performance: Bandwidth for DirectX 10/Vista
  - Power efficiency and battery life (avoid the local frame buffer)
- Power efficiency limitations:
  - Frequency and voltage cross dependencies limit the granularity of power savings
  - Frequency agility

To be addressed without changing the CPU core!

Addressing these Challenges

- Increase system bandwidth for UMA solutions
  - Performance: HyperTransport™ 3, dedicated display refresh virtual channel, maximize DRAM efficiency
  - Power: HT3 power management extensions, Memory controller on its own power plane
- Power efficiency
  - Split power planes, independent frequency selection (core0, core1, NB are independent)
  - Fast frequency changes without PLL relock using digital frequency synthesizers
  - Improved efficiency allows lower power for fixed workloads, or higher performance for fixed power

Introducing "Griffin"

Enhanced performance and battery life

New infrastructure features optimized for the mobile segment

System on chip (SOC) methodology to accelerate design and integration
- New mobile optimized Northbridge
- Power optimized DDR2
- Enhanced HyperTransport™ 3 connectivity
- Integrating 65nm cores and larger L2 caches

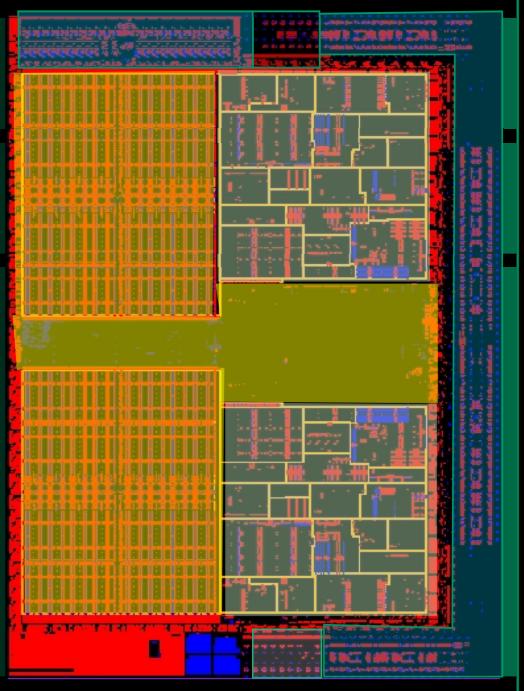

A Few Details

Dual AMD64 CPU cores

Dedicated 1 MB L2 caches

Dual channel DDR2 interface
- 64 bits per channel
- All speeds supported up to DDR2-800
- 12.8 GB/s peak memory bandwidth
- 16 GB max memory configuration

HyperTransport™ 3 I/O link
- Speeds supported up to 2.6 GHz
- 16x16 bit
- 10.4 GB/s simultaneous peak bandwidth in each direction
- Dynamic link power management

160 mm², 225.6 million transistors

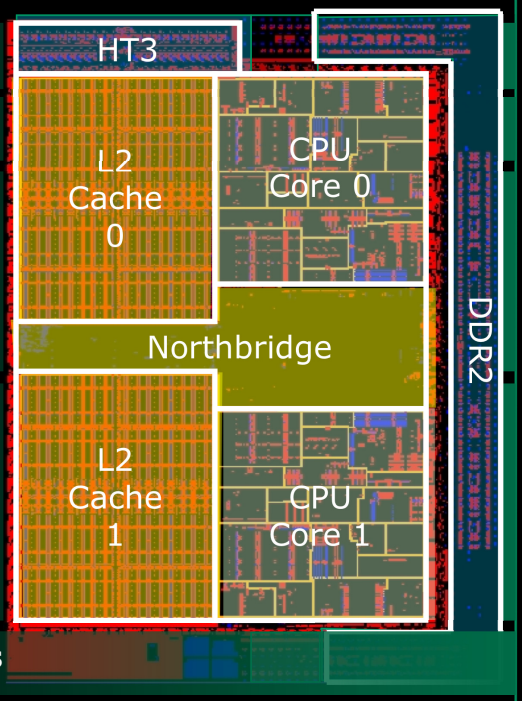

Mobile Optimized Memory Controller

New integrated memory controller features:
- Improvements in DRAM efficiency
- Improved DRAM prefetcher
- Dedicated Display Refresh channel

Extended battery life with new power saving features
- Operates on separate power plane, and at a lower voltage than the cores
- Enables C4 Deeper Sleep on systems with UMA graphics without the need for local frame buffer

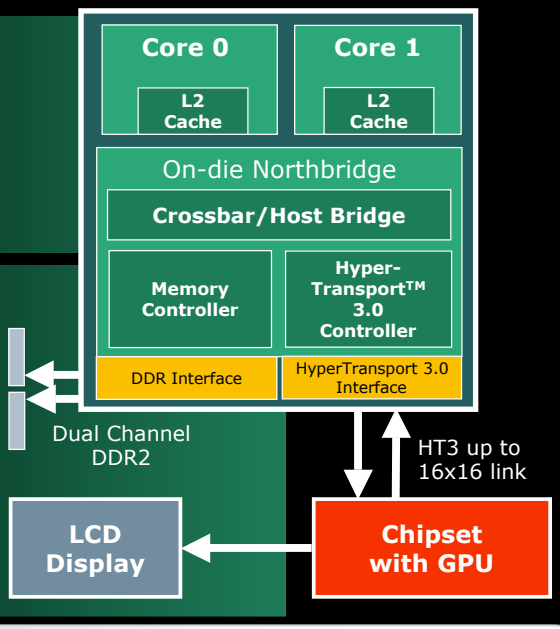

Cost and power efficient solution for UMA

DRAM Efficiency Maximization

AMD Smarter Choice

Aggressive use of bypass paths to minimize idle latency

Bank state tracking

- Chip selects are interleaved to increase number of distinct banks
- Unganging of DRAM channels further increases number of banks
- 16 banks per channel are tracked, using LRU algorithm
- Pages can be closed dynamically, based on bank access pattern

Out of Order (OOO) scheduling of requests, based on:

- Request priority (programmable by type: low/med/high)
- Page status (miss/hit/conflict)

Write bursting

- Writes are accumulated and done at once to minimize bus turnaround

Advanced DRAM Prefetcher

AMD Smarter Choice

Memory Prefetch Table (MPT) has entries for 8 outstanding threads

Prefetches 2 requests ahead

Capable of tracking single (+n, +n, +n) or double (+n, +m, +n) strided patterns, with strides up to ±4 lines

When request is received:

- *stride2* = change in address + *stride1*
- *stride1* = change in address

stride2 + last is predicted address; prefetches are issued when confidence threshold achieved

MPT

0	*stride1*	*stride2*	*addr*	*count*
1	*stride1*	*stride2*	*addr*	*count*
2	*stride1*	*stride2*	*addr*	*count*
3	*stride1*	*stride2*	*addr*	*count*
4	*stride1*	*stride2*	*addr*	*count*
5	*stride1*	*stride2*	*addr*	*count*
6	*stride1*	*stride2*	*addr*	*count*
7	*stride1*	*stride2*	*addr*	*count*

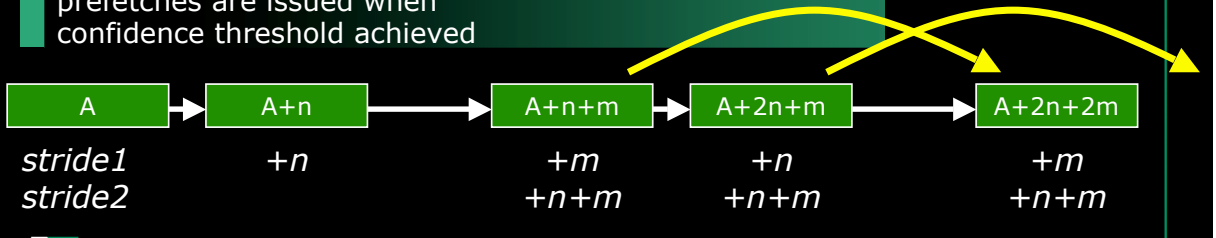

Display Refresh Optimization

AMD Smarter Choice

High bandwidth, high latency, but latency guarantee required to avoid display buffer underrun

Doesn't fit well in existing HyperTransport™ VC sets

- Base channels provide no latency guarantees
- Isoc channels are low bandwidth; can starve other traffic

Requests arrive via HyperTransport™ Isoc channel

- Detected by decoding coherence and ordering requirements
- Has dedicated buffer and routing resources internally
- Chipset must manage interaction with Isoc traffic

Memory controller priority is variable based on age

Dynamic Performance Scaling Capabilities

AMD Smarter Choice

Reduced power consumption to increase battery life

- Separate voltage planes for each core
- Each core can operate at independent frequency and voltage

Increased performance with instantaneous frequency transitioning

- No PLL relock required
- Minimize power consumption by always running at optimal P-state
- Lower operating minimum p-state
- Reduced processor utilization with simplified p-state transitions

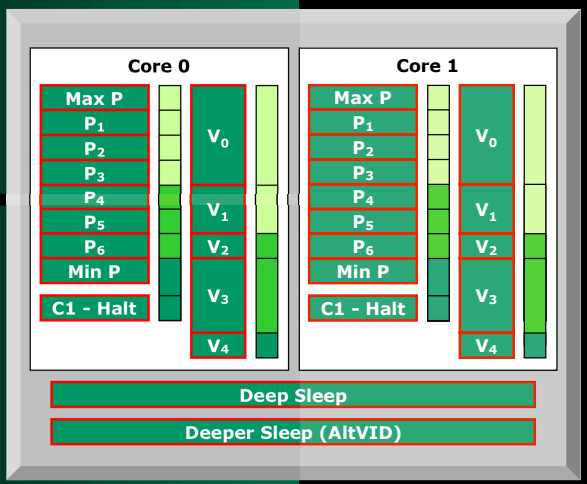

Increased battery life with advanced power management features

Power-optimized HyperTransport™ 3

AMD Smarter Choice

Increased performance with over 3x increase in peak I/O bandwidth

- Delivers increased bandwidth for DX10, a requirement for 2008 Windows Vista™ Premium Logo

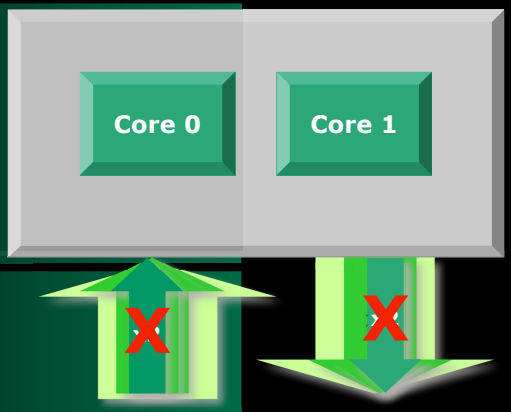

Extended battery life with power reduction features

- Dynamic scaling of link widths
- Disconnecting HyperTransport™ when idle, even while cores are executing
- HW-autonomous, no SW support needed

HyperTransport™ 3 increases performance while helping to extend battery life with power management features

Note: Animated scaling of link widths shown here are for illustration purposes. HT disconnect is required between each link width change. Actual link width scaling may vary based on implementation.

Voltage Planes and Control

AMD Smarter Choice

Fixed Analog and I/O Voltage Planes

- VDDIO & VTT for DDR PHY
- VDDA for on-die PLL
- VLDT for HyperTransport PHY

Independently-Variable Voltage Planes

- VDD0 for CPU core 0
- VDD1 for CPU core 1
- VDDNB for on-die NorthBridge

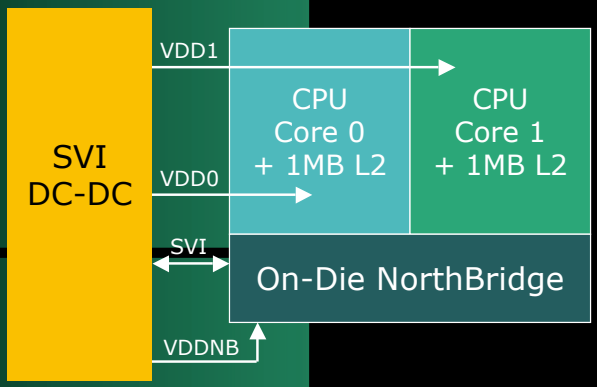

Serial VID Interface (SVI) Protocol provides pin-efficient means to control multiple independent voltage planes

Fine Grain Power Management

AMD Smarter Choice

CPU core power-state transitions

- Simple software interface
- Cores continue execution while frequency changes are in progress

Autonomous hardware power management

- CPU and chipset work together to establish the most power efficient settings
- Eliminates reliance on software for maximum power efficiency
- The CPU informs the chipset when P-states change or HALT condition reached
- The chipset monitors I/O traffic and CPU state and establishes the optimal power management profile for a given set of system conditions

Current-generation power management schemes remain supported

Power efficiency via dynamic hardware power management

Autonomous Hardware Power Management

AMD Smarter Choice

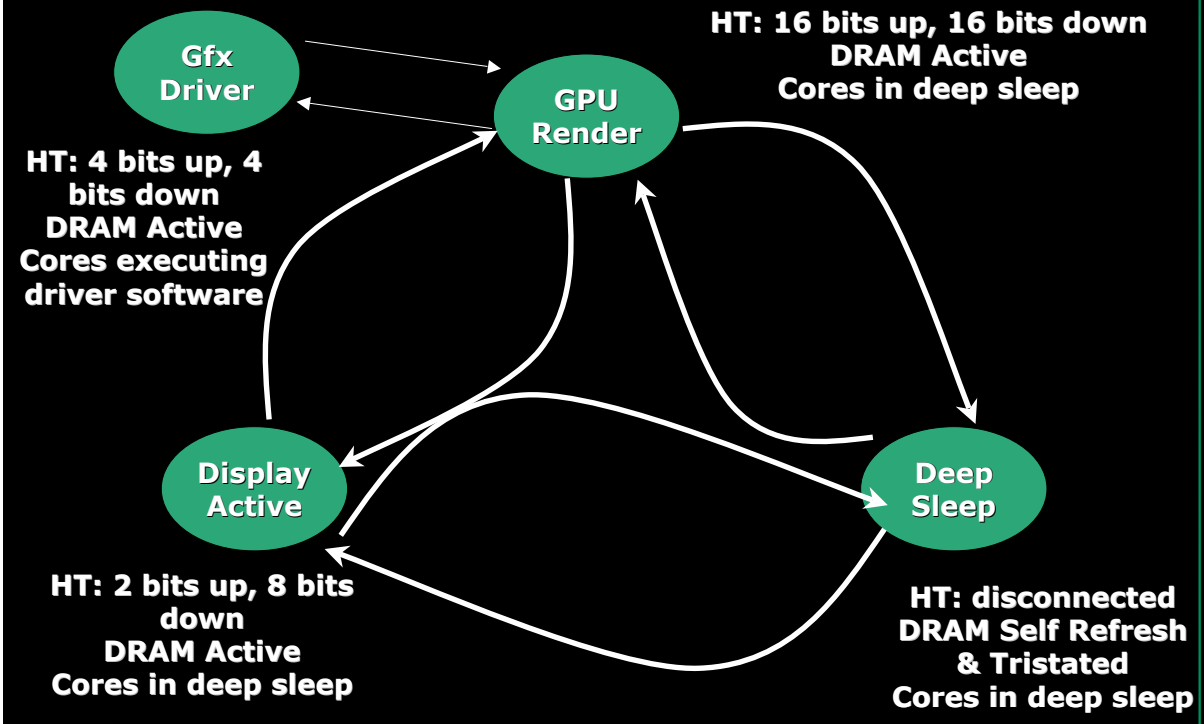

Power and Performance Tradeoffs

Feature	Attributes
Independent CPU core voltage planes and frequency selection	Power consumption matches CPU performance delivered
Separate voltage plane for on-die NorthBridge	Enables CPU deep sleep with integrated graphics.
Dynamic HT ™ link power management	Power consumption matches interconnect bandwidth delivered
Autonomous hardware control of CPU core deep sleep state	Increased residency in CPU deep sleep state
Autonomous hardware control of DRAM self-refresh state	Increased residency in DRAM sleep state
CPU core deep sleep wakeup to service probes at lowest P-State	Increased residency in CPU deep sleep state

Griffin Power Management Visualization

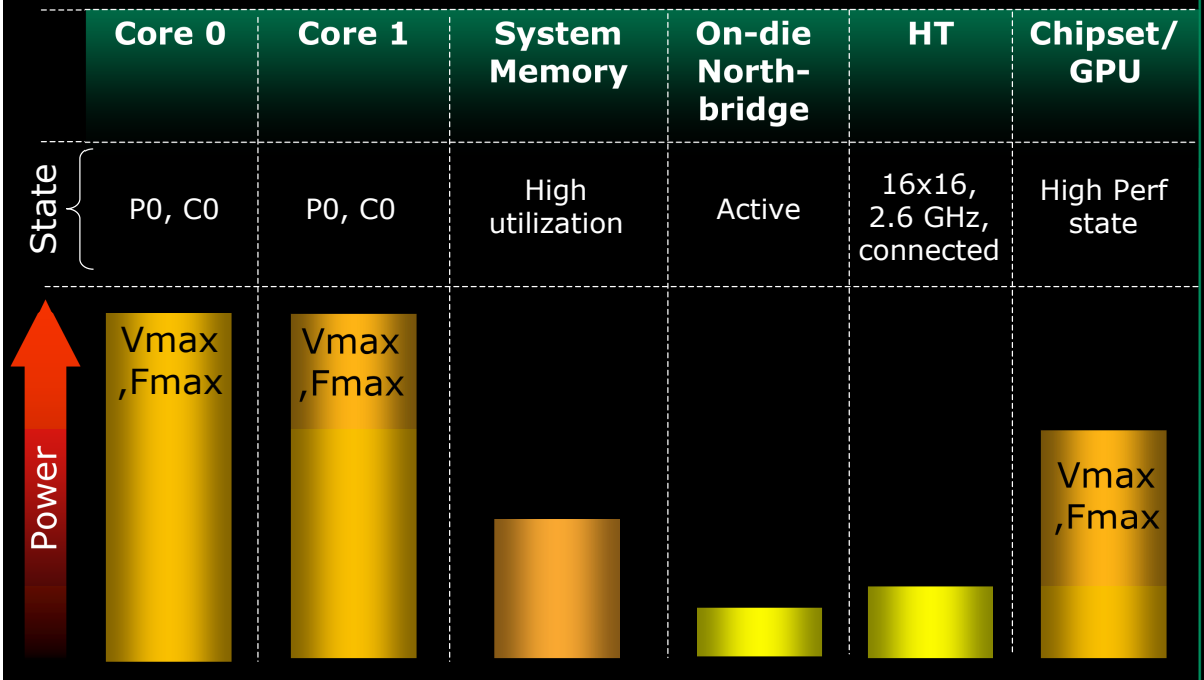

Griffin Power Management Visualization

AMD Smarter Choice

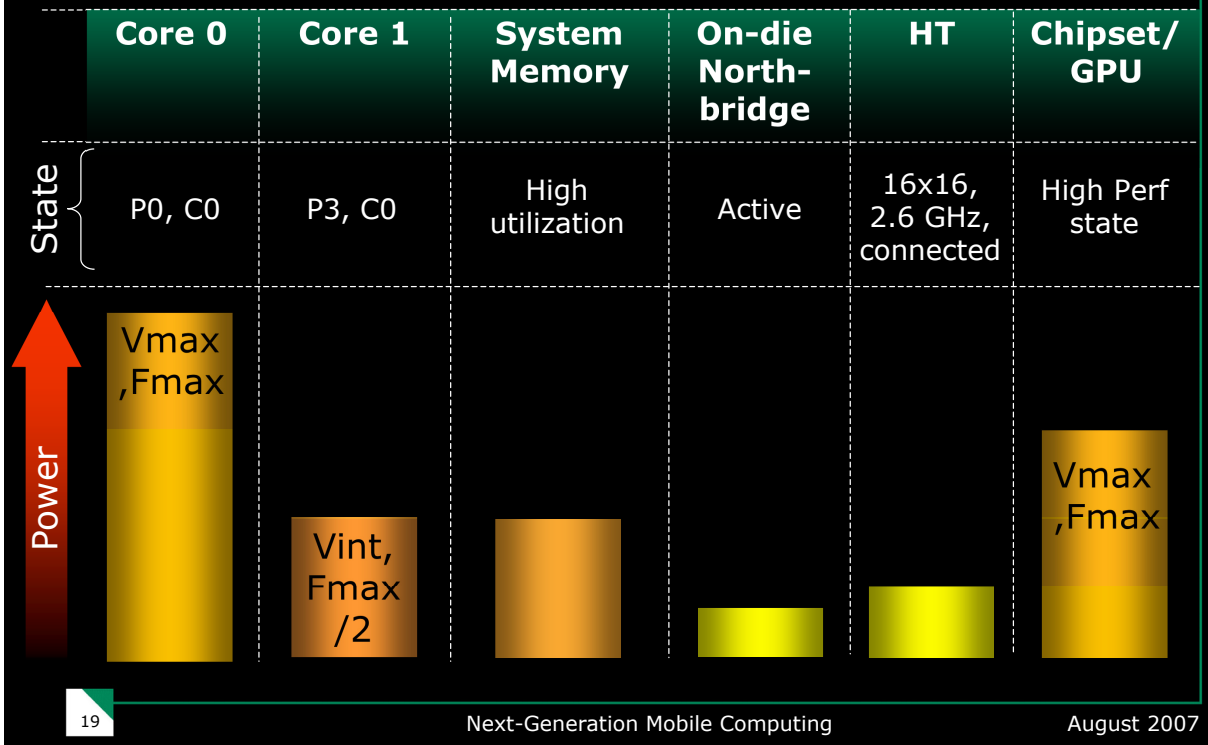

Griffin Power Management Visualization

AMD Smarter Choice

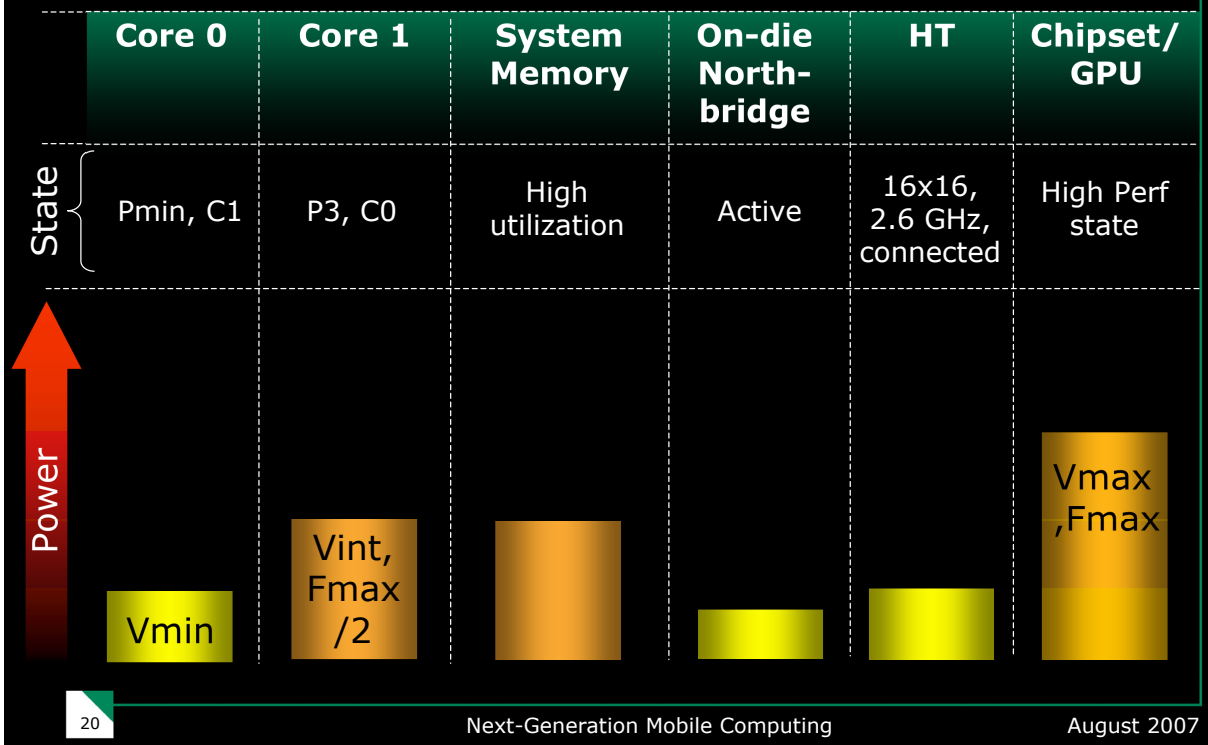

Griffin Power Management Visualization

AMD Smarter Choice

	Core 0	Core 1	System Memory	On-die North-bridge	HT	Chipset/ GPU
State	Pmin, C1	Pmin, C1	Self refresh	Clocks gated	Dis-connected	Low Perf state

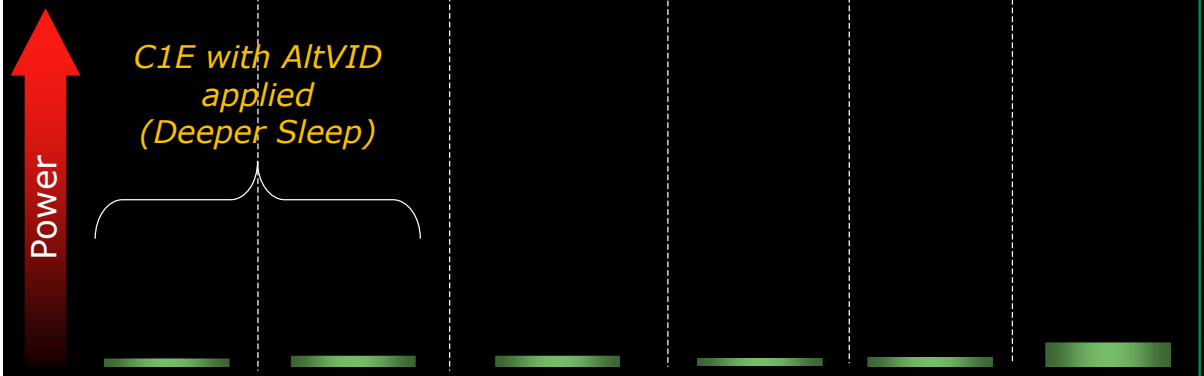

Griffin Power Management Visualization

AMD Smarter Choice

	Core 0	Core 1	System Memory	On-die North-bridge	HT	Chipset/ GPU
State	Pmin, C1	Pmin, C1	Active	Active	2x8, 2.6 GHz, connected	Low Perf state (refresh)

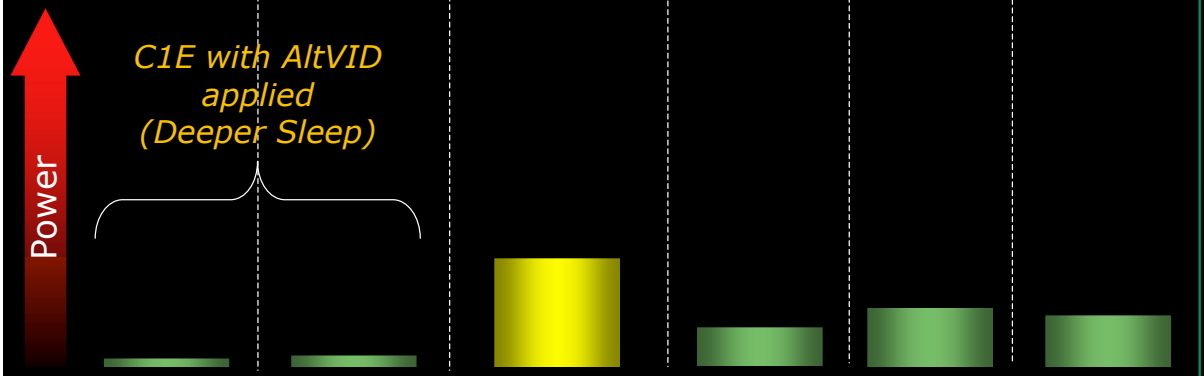

AMD Multi-Point Thermal Control

AMD Smarter Choice

Multiple on-die thermal sensors

Simplified thermal management

- Designed to automatically reduce p-state when temperature exceeds pre-defined limit

New memory power management

- External thermal monitor can force reduction in memory temperature with MEMHOT signal

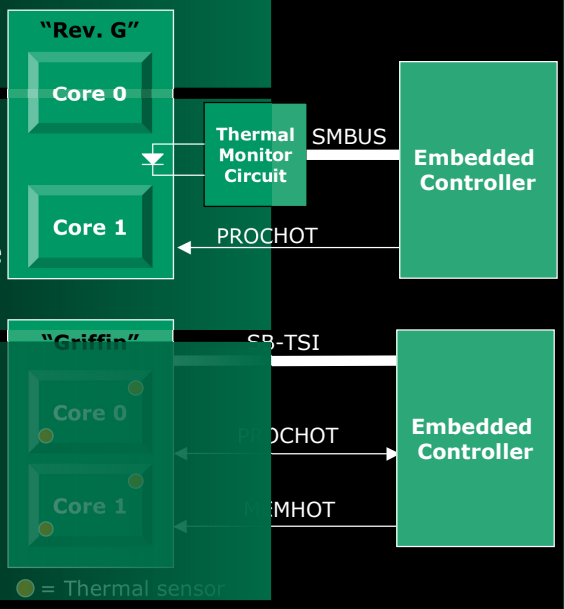

Simple and accurate thermal management

Summary

AMD Smarter Choice

"Griffin" is an important evolutionary step in AMD's notebook product portfolio

- Greater memory efficiency for improved processor and graphics bandwidth
- Enhanced power management
- Improved thermal management

Optimized for UMA power/performance while maintaining direct CPU-Memory connection

- Simplified, reliable platform architecture
- Low platform cost
- Maximal CPU performance

Tighter integration of the chipset and CPU an important milestone to Fusion

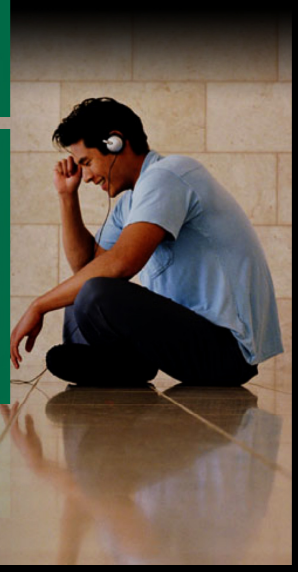

Disclaimer

The information presented in this document is for information purposes only.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including, but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN OR FOR THE PERFORMANCE OR OPERATION OF ANY PERSON, INCLUDING, WITHOUT LIMITATION, ANY LOST PROFITS, BUSINESS INTERRUPTION, DAMAGE TO OR DESTRUCTION OF PROPERTY, OR LOSS OF PROGRAMS OR OTHER DATA, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Trademark Attribution

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG. HyperTransport is a licensed trademark of the HyperTransport Consortium. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.